

Table of Contents

Week of 100531.....	1
Daily WLCG Operations Call details.....	1
WLCG Service Incidents, Interventions and Availability, Change / Risk Assessments.....	1
General Information.....	1
Monday:.....	1
Tuesday:.....	3
Wednesday.....	6
Thursday.....	8
Friday.....	10

Week of 100531

Daily WLCG Operations Call details

To join the call, at 15.00 CE(S)T Monday to Friday inclusive (in CERN 513 R-068) do one of the following:

1. Dial +41227676000 (Main) and enter access code 0119168, or
2. To have the system call you, [click here](#)
3. The scod rota for the next few weeks is at [ScodRota](#)

WLCG Service Incidents, Interventions and Availability, Change / Risk Assessments

VO Summaries of Site Usability				SIRs & Broadcasts		Change assessments
ALICE	ATLAS	CMS	LHCb	WLCG Service Incident Reports	Broadcast archive	CASTOR Change Assessments

General Information

General Information				GGUS Information	LHC Machine Information
CERN IT status board	M/W PPSCoordinationWorkLog	WLCG Baseline Versions	WLCG Blogs Blogs	GgusInformation	Sharepoint site site - Cooldown Status Status - News News

Monday:

Attendance: local(Maria, Jamie, Jean-Philippe, Simone, Harry, Ignacio, Jan, Eva, Carlos, IanF, Patricia, Nilo, MariaDZ, Manuel);remote(Jon, Angela, Joel, Tore Mausset (NDGF), Alexander Verkooijen (NL-T1), Gang, Gonzalo).

Experiments round table:

- ATLAS reports -
 - ◆ Tails of reprocessing data distribution being drained. Backlog only in the BNL-LYON channel: per file transfer rate down to 2 to 3 MB/s, not clear whether the problem is in LYON, BNL or somewhere in the middle. Sites contacted - GGUS:58646 [GGUS:58646](#) assigned to Lyon (since BNL is on holiday).
 - ◆ Alarm ticket against CERN on Friday night because of AFS being very slow, affecting T0 operations. The problem was that some ATLAS SW releases were not replicated to the RO volume, therefore all analysis jobs were hitting the RW volume. Problem fixed by ATLAS release experts at around midnight.
 - ◆ FTS 2.2.4 in PPS looks OK for the moment (no stale jobs) - upgrade possibly tomorrow? What do other experiments think...
 - ◆ 30' of SRM ATLAS red - monitoring problem? Yes - artifact of other problems
- CMS reports -
 - ◆ T0 Highlights
 - ◇ No collision running over the weekend. An investigation for later Heavy Ion running

was conducted on Sunday and there will be a second test tomorrow. In CMS heavy ion events are much larger because the tracker zero suppression is not done in hardware. (>10MB/event) - rates into CASTOR 1.6GB/s - schedule test end of summer with other experiments. No problem for FTS upgrade tomorrow.

- ◆ T1 Highlights
 - ◇ Reprocessing with 3_6_1 is mostly finished. Skims are in progress on the reprocessed data.
- ◆ T2 Highlights
 - ◇ MC production as usual
- ◆ Tier-1 Tickets: #114800: Batch Priority Policy at PIC CMS Pilot jobs are not getting sufficient batch slots to finish the reprocessing #114751: Lost MC data at RAL Bad tape at RAL list 16 MC files: Gonzalo (PIC) - didn't see any GGUS. IanF - using Savannah - GGUS bridge...Will put in GGUS manually.
- ALICE reports [↗](#) - TECHNICAL STOP: No reconstruction activities and no MC production. There are two analysis trains running at this moment. In terms of transfers, about 11 TB have been transferred during the last weekend with an average speed of 32MB/s. ALICE has been informed about the alarms exercise foreseen for this week. Efficiency of transfers to FZK? Due to attempt to transfer those files that have been lost.
 - ◆ T0 site
 - ◇ Intervention announced on Friday for today: there was a bad router affecting 11 disk servers in castoralice/alicedisk and 2 disk servers in t0alice with network errors and disconnects. The network people asked when could an intervention be scheduled to sort out this problem. Online and Offline representative have agreed with the network team to have this operation today at 9:30. At around 10:00 the team announced the operation had successfully concluded
 - ◇ No issues to report in terms of the T0 resources during the weekend
 - ◇ Ticket: GGUS:58587 [↗](#) submitted on Friday to the T0, still waiting for actions
 - ◇ LanDB set defined for ALICE and reported last week still appears to be empty. Manuel: created automatically in CDB but connection to LANDB not really working.. Recipe for HTAR?
 - ◆ T1 sites
 - ◇ Low activity at the T1 sites at this moment
 - ◇ FZK: ALICE expert has announced the setup of the CREAM1.6 service at the T1. The service has been put in production
 - ◆ T2 sites
 - ◇ Ticket GGUS:58590 [↗](#) submitted to Clermont: SOLVED
 - ◇ Ticket GGUS:58591 [↗](#) submitted to Cagliari: Ongoing
 - ◆ KIT - one CREAM CE with 1.6 - other will be upgraded in the next month.
- LHCb reports -
 - ◆ Experiment activities:
 - ◆ MC production ongoing. Plan some reconstruction mid-week once new LHCb stack released.
 - ◆ GGUS (or RT) tickets:
 - ◇ T0: 1
 - ◇ T1: 0
 - ◇ T2: 4
 - ◆ Issues at the sites and services
 - ◇ T0 site issues:
 - problem with LHCbMDST fixed on Friday by adding 11 more disk servers into the service class serving this space token. During the week end no more problem observed or reported by users.
 - AFS problem starting from 10:30. (ALARM GGUS:58643 [↗](#)). Many users affected, many monitoring probes submitting jobs via acrontab affected,

e-logbook not working and other services relying on were severely affected. Shared area for grid jobs not affected however. Issue boiled down to a power failure on a part of the rack. Ignacio - problem was a switch that was no longer responding. AFS team preparing a report of what happened. Classify this as a network problem, rather than an AFS problem. Joel - took a long time before SSB was updated to reflect this problem. Ignacio - first entry in timeline 10:29. Resolved ~11:15.

◇ T1 site issues:

- Issue at Lyon throttling activities last week was not due to a limitation of the new very-long queue but to CREAMCE jobs status wrongly reported by the gLite WMS.
- Issue at SARA with SAM tests. It seems due to the CERN CA expired in one gridftp server (GGUS:58647 [↗](#))

◇ T2 sites issues:

- Shared area issues at CBPF and PSNC. Jobs failing at UKI-LT2-QMUL and UFRJ

Sites / Services round table:

- FNAL - ntr
- KIT - problems on w/e with one CE. gridftp server died. Back to normal now. CE3.
- NDGF - ntr
- ASGC - last Friday mentioned some errors at ASGC. Quickly verified - AFS repaired and problem fixed. File had become a directory(?)
- NL-T1: 1) Upgraded network between storage-compute cluster to 160Gbps. 2) GGUS:58647 [↗](#) (as under LHCb report)
- PIC - tomorrow due to LHC tech stop at 12:00 have scheduled 1h intervention (expect to last 10') to deploy tape protection in dCache. Should be transparent. CMS issue - will follow offline. CMS production role has fair share of 45%. About 20 jobs with this DN arrive to PIC per day. Follow through ticket.
- CERN DB - on Friday some problems on ATLAS offline DB caused by Panda and PVSS applications - high load. Fixed on Friday. This morning during patching of CMS online had some problems with Quattor profiles and thus intervention took longer than expected. Had to kill all instances as were blocking each other. Progressing with rest of patches for other experiments.
- CERN CASTOR PUBLIC - problem of 90' during morning - routine operation restarting LSF. Still investigating. Service is back but investigating why it went wrong. Affected monitoring.

AOB:

Tuesday:

Attendance: local(Gavin, Jean-Philippe, Ignacio, Simone, Harry, Jamie, Maria, Kate, Manuel, Steve, Tim, Patricia, Dirk, MariaDZ, Nilo, Lola);remote(Michael, Jon, Angela, Joel, Gonzalo, Ronald, Gang, Jeremy, Tiju, Stefano Zani (INFN TIER1), Rob, IanF).

Experiments round table:

- ATLAS reports -
 - ◆ This morning the ATLAS DDM dashboard was unavailable for approx 1h. The problem was observed at 10:45, during the rolling intervention to the database. Trying to restart the dashboard apache server as in the procedure did not work, so dashboard support has been contacted.

Tuesday:

- ◊ Explanation from David Tucket: "For some time we could not make a connection to atlas_dashboard database. Probably due to the intervention you indicated (although it was only supposed to affect existing sessions). The database is now available and I have restarted the DDM Dashboard, consumers and agents. All seems to be running normally. I'll continue to monitor the situation and look into why the various systems did not recover themselves when the database became available again."
 - ◊ As far as I can tell, no other ATLAS service was affected.
 - ◆ After lunch (report from Florbella at 14:15) the all ATLR RAC become unaccessible.
 - ◊ Report from Florbella "All services on ATLR are down due to a cluster problem related with the rolling security updates. All services mentioned (DQ2 Services , Panda, Conditions,T0, Prodsys) as well as all related offline applications are down. An update will be posted when services are restored."
 - ◊ In addition, also the Dashboard is down again. Is it the same problem? I thought Dashboard was on WLCG RAC as well as ATLAS LFC and the latter is OK.
 - ◆ Alarms for the ATLAS CERN critical tests have all been sent (11 tests) * GGUS:58654[↗](#), GGUS:58656[↗](#), GGUS:58657[↗](#), GGUS:58674[↗](#), GGUS:58675[↗](#), GGUS:58676[↗](#), GGUS:58677[↗](#), GGUS:58687[↗](#), GGUS:58688[↗](#), GGUS:58689[↗](#), GGUS:58690[↗](#)
 - ◆ BNL-LYON slow transfers issue still being investigated (GGUS:58646[↗](#))
 - ◆ FTS 2.2.4 - still see no problems after 4 days of running
- CMS reports -
 - ◆ T0 Highlights
 - ◊ Another Heavy Ion Test Today
 - ◆ T1 Highlights
 - ◊ Mopping up 3_6_1 reco. Skims are in progress on the reprocessed data.
 - ◆ T2 Highlights
 - ◊ MC production as usual
 - ◆ Oracle RAC node still down
 - ALICE reports[↗](#) - GENERAL INFORMATION: No reconstruction nor MC activities. Two User analysis trains currently in production
Raw data transfer activities from the T0 to the T1 sites ongoing to CNAF, FZK, NDGF and RAL
 - ◆ T0 site
 - ◊ LanDB set issue reported yesterday (still the set appeared empty): SOLVED. This morning the solution has been confirmed by PES experts. The next step is to inform the security team. They will have to change the firewall exceptions from the previous LanDB to the current one. Once this operation is done, the previous LanDB can be deprecated.
 - ◊ ALARM ticket tests at the T0: 4 alarms already submitted for CASTOR, xrootd, VOBOXES and CREAM. Tickets for the VOBOX and for CREAM have been closed. CASTOR experts reported the piquet had not been called, so something in the chain did not work as expected.
 - ◆ T1 sites
 - ◊ FZK. Feedback about CREAM1.6: the site has been configured in LDAP to use exclusively the CREAM1.6 system at the site. The system had been able to run more than 1000 concurrent jobs during the last night. No issues to report
 - ◆ T2 sites
 - ◊ Ticket GGUS:58591[↗](#) submitted to Cagliari: involving the ALICE experts in Italy to solve the issue
 - ◊ Minor operations required this morning in Madrid. Restart of all local services
 - LHCb reports -
 - ◆ Experiment activities: (see full report)
 - ◆ MC production ongoing

- ◆ Tested and verified the GGUS ALARM workflow at CERN for the LHCb critical services
- ◆ GGUS (or RT) tickets:
 - ◇ T0: 0
 - ◇ T1: 1
 - ◇ T2: 2
- ◆ Issues at the sites and services
 - ◇ T0 site issues:
 - none
 - ◇ T1 site issues:
 - Issue at SARA with SAM tests. Confirmed that it was one diskserver with CERN CA expired. Some other failures observed due to one of the diskserver in maintenance (GGUS:58647 [↗](#)).
 - dCache sites in general: (Input for the T1 coordination meeting). Details available on GGUS:58650 [↗](#)
 - ◇ T2 sites issues:
 - GRISU-SPACI-LECCE: shared area issues
 - UK-T2 sites upload issue: GGUS open for middleware: GGUS:58605 [↗](#)
Jeremy - not much to comment at moment. Have issues at several sites. NAT sites + gridftp but ok at some sites. Still investigating.

Sites / Services round table:

- BNL - big maintenance until 18:00 CERN time. Most services affected. Other things regarding network connectivity for storage servers - some will experience small outage when moving to a new switch - not expected to affect operations
- FNAL - ntr
- KIT - ntr
- ASGC - ARGUS become production service yesterday
- PIC - scheduled intervention at 12:00 went ok - no impact. Upgraded dCache to new subversion which incorporates tape protection. Ticket from CMS: to clarify - are you waiting for site input? IanF - waiting for confirmation from people who developed pilot information - not waiting for anything from PIC. Tested in manual way what should be done by pilot factory. Slow but successful. Need some answers from developers who were on holiday yesterday.
- NL-T1 - SARA found dCache bug when transfer failing or aborted. dCache DB keeps some filespace records that last a couple of hours. New transfers to same destination will fail. Known bug or not? Checking. Some FTS transfer failures.
- RAL - at risk for UPS test. Went ok.
- CNAF - ntr
- GridPP - ntr
- OSG - couple of minor transfer issues at BNL reported this w/e - nothing of large impact. Quiet w/e.

- CERN CMS Castor instance was unable to record new files between 1am and 5am this morning. Following a piquet call and developer assistance, the problem has been resolved and the instance is running smoothly again. Jobs may have failed during this period with timeouts. Tim - probably reported by operator around 1am. Writes were blocking - brought in developers. Anyone trying to write into CASTOR at this time would have been blocked. Will introduce a fix and deploy around systems. Relatively unusual situation with garbage collection at same time as disk-disk copy. A short-term fix and something in next release. Detailed in full report that will come.

- CERN myproxy, Tomorrow morning a lot of old expired proxies dating back to 2008 will be deleted from myproxy.cern.ch. This is not expected to cause any problems. A disk copy of them will be kept for a time in case problems are reported.

- CERN FTS 2.2.4 upgrade - ATLAS gives green light (pending on LCGR). CMS - in favour of

upgrade. No experts around early this week so no test experience. As can be rolled-back out ok with upgrade.

- CERN DB: 2 interventions scheduled for today, LCGR: should have been rolling. Shutdown of one instance - m/c rebooted and kernel panic. Turned out that problem from yesterday with CMS repeated and clusterware went crazy. Required downtime and reboot of instances. Problem repeated with ATLAS. Thus DB taken offline and not rolling. Maria - would like to see full understanding of problem before patching continues. Clearly correlation between applying patch and carrying on. Also for external sites? Also open question on CMS node - vendor has 12h to react and replace memory. Vendor says until Friday.. Pushing support. Need to keep spare nodes? (LHCb patching currently ongoing...)

AOB:

Wednesday

Attendance: local(Jean-Philippe, Stephen, Jamie, Maria, Eva, Nilo, Gavin, Patricia,Roberto);remote(Jon, Xavier, Rolf, Gang, Onno, Jeremy, Tiju).

Experiments round table:

- ATLAS reports -
 - ◆ After the installation of the Oracle patch yesterday, there have been major instabilities at various ATLAS services. The most affected seems to be prodsys (major impact on group NTUPLE production to be done by the end of the week). ATONR patch has been rolled back at 1PM, ATLR rolled back at 14:00.
 - ◆ DDM Site Services being migrated to SLC5 starting from today, for the next couple of days. No service interruption foreseen, possibly a few "FILE EXIST" errors and some SLS artifact during the migration period.
 - ◆ DB - problem seems to be related to use of COOL and auditing. Rolled back ATLAS, rolling back LHCb. Other sites should roll back.
- CMS reports -
 - ◆ T0 Highlights
 - ◇ Point5->T0 transfers recovering from upgrade yesterday
 - ◆ T1 Highlights
 - ◇ Mopping up 3_6_1 reco. Skims tailing off on the reprocessed data.
 - ◇ Snapshot of ORCOFF (conditions) tomorrow afternoon, possible it could effect T1 reprocessing
 - ◆ T2 Highlights
 - ◇ MC production as usual
 - ◆ Oracle RAC node still down (after almost one week.) See https://twiki.cern.ch/twiki/bin/view/PDBService/DBServicePostMortem#CMSR_node_broken_May_2
- ALICE reports [↗](#) - GENERAL INFORMATION: Reconstruction activities restarted. Transfers: Considering the rates defined per site and based on the available resources provided for ALICE, all T1 sites have now achieved their part of raw data but CNAF and NDGF. Raw data transfers priority has been give today to these two sites.
 - ◆ T0 site
 - ◇ None of the local CREAM-CEs at CERN is working. Top priority GGUS ticket: 58718
 - ◇ Ticket 58587 submitted last week and concerning to missing c++ headers in 4 WNs. It seems the nodes have been taken out of production for further investigations. Any news about it?

- ◊ Security team has been informed about the population of the IT CC CAF NODES LanDB set.
- ◆ T1 sites
 - ◊ FZK: Yesterday the site admins informed about an intermittent problem observed in SAM: From time to time the SAM tests announce that the proxy registration procedure into one of the local VOBOXES fails. Cheking the log files of the SAM UIs we found that the VOBOX test suite had not properly run at FZK. It was manually executed and this time it worked.
 - ◊ CCIN2P3: The ALICE expert informed yerteday about the wrong results reported by MonaLisa concerning the status of the VOBOXES at this T1 site. However the site was performing correctly. The corresponding test suite used by MonaLisa is having problems with the shell defined in the VOBOXES at this site and therefore the problem is site independent. To be discussed during the next ALICE TF meeting this week.
- ◆ T2 sites
 - ◊ Ticket 58591 to Cagliari: Escalated
 - ◊ Ticket 58719 to GRIF_IRFU submitted this morning: The local CREAM-CE provides timeout errors at submission time. Site out of production due to this problem.
- LHCb reports -
 - ◆ Experiment activities:
 - ◆ MC production ongoing
 - ◆ GGUS (or RT) tickets:
 - ◊ T0: 0
 - ◊ T1: 0
 - ◊ T2: 0
 - ◆ Issues at the sites and services
 - ◊ T0 site issues:
 - Load on lhcbmdst despite the new 11 disk servers added on Friday. This because data reside exclusively on the old servers and load is not spread. Moved to xrootd that does not trigger a LSF job for reading files already staged in disk
 - ◊ T1 site issues:
 - IN2p3: jobs killed because exceeding memory limit (2GB) while they are expected to consume 1.5GB (as per VO Id Card). May be a memory leak on DaVinci application.
 - ◊ T2 sites issues:
 - T2 UK sites upload issue: this has to be escalated at the T1 coordination meeting being now a long standing problem and has to be address systematically in a coordinated way.

Sites / Services round table:

- FNAL - ntr
- KIT - ntr
- IN2P3 - ntr Q (Simone) - progress on understanding traffic BNL-IN2P3. A: see Monday: GGUS:58646 [↗](#) assigned to Lyon (there is no update - Ed.)
- ASGC - ntr
- NL-T1 - ntr
- RAL - at risk for patching Oracle DBs. Not completed. Extended downtime.
- GridPP - ntr
- OSG - Item for BNL ATLAS GGUS:58628 [↗](#)

- CERN FTS, Upgrade of prod services to 2.2.4 completed with out incident. Intervention started at 10:00 CEST and completed at 10:36 for T0 service and 10:52 for T2 service. During intervention new work was accepted but not started till after intervention. Q: is recommendation for T1s to do this? Discuss at T1SCM tomorrow.
- CERN DB Postmortem on "rolling" upgrade. Which DBs being patched at RAL? A: FTS & ATLAS LFC.

AOB:

- GGUS test alarms - none from CMS! Critical services common - would be good to have at least one test from CMS to check that alarm goes to CMS experiment. See the extract from last Thursday: A round of the Periodic ALARM tests full chain will take place next week Mon-Wed. The Tier0 ONLY is concerned this time. Steps to follow are in <https://savannah.cern.ch/support/?114705> and services to test are in <https://twiki.cern.ch/twiki/bin/view/LCG/WLCGCriticalServices>

Thursday

Attendance: local(Jean-Philippe, Patricia, Roberto, Manuel, Harry, Jamie, Maria, Simone, MariaDZ, Dirk, Lola, Eva, Maarten, Nilo, Kate, Julia, Stephane);remote(cristina vistoli, Jon, Stephen, Gang, Joel, Rolf, Ronald, Gonzalo, Jeremy, Rob, Gareth).

Experiments round table:

- ATLAS reports -
 - ◆ ATLAS DBs at CERN now stable. Last interruption was yesterday at approx 4PM when ATLR was rebooted. Not clear to ATLAS if it was part of the rollback or side effect or some other issue. [Eva - problem was caused by high load on DB causing swap and one node got local filesystem corrupted. Had to fix to restart this. High load on remaining RAC caused problems. Simone - should I involve people from ATLAS applications side? Eva - Luca has contacted application people. Marcin following up.]
 - ◆ GGUS:58646 (slow transfers BNL to LYON) is now high priority. Problem persists after 4 days from notification and the activity will ramp up in the next days. I added more infos to the ticket showing that BNL->CERN single file transfer rate is OK (13 to 40 MB/s) and that therefore the issue is not in BNL. (Tried same transfers CERN-Lyon and problem also there. BNL-CERN single file transfer ok - 30-40MB/s. Same issue CERN-Lyon means as soon as raw data comes 4MB/s will not be enough in case of big dataset & long run of ~24h. (For exports to Lyon...)) [Rob - OSG - never opened against BNL. Don't need to open against BNL it seems? A: right - ticket opened against Lyon due to US holiday - seemed to be correct. BNL in cc: in original thread and helped with e.g. iperf tests.]
- CMS reports -
 - ◆ T0 Highlights
 - ◇ Point5->T0 transfers still recovering from upgrade Tuesday
 - ◆ T1 Highlights
 - ◇ PIC FTS problems: Savannah GGUS against developers.[Gonzalo = opened on PIC Mon/Tue issues with glide-ins to PIC. Following up. Related GGUS ticket opened yesterday with no more info. Are you expecting more from PIC? MariaDZ - pursued IanF's q on Monday as why Savannah-GGUS bridge didn't work. A mandatory field (Status field was not set manually to "GGUS on-hold) was not completed - lack of value in this field caused the problem. Done manually. Stephen - will find out]
 - ◇ Skims tailing off on the reprocessed data.
 - ◇ Snapshot of ORCOFF this afternoon, possible it could effect T1 reprocessing

- ◆ T2 Highlights
 - ◇ MC production as usual
- ◆ Oracle RAC node back in production 7pm last night.
- ◆ Test GGUS Alarm ticket opened yesterday after meeting: GGUS Ticket. Yesterday I was not in /cms/ALARM but added last night. Can still not open ALARM tickets, is the problem reported by me here in November still the case? Updates from VOMS groups into GGUS automatic? MariaDZ - happen once per night. Please open GGUS ticket as public holiday [☞](#) in Germany today.
- ALICE reports [☞](#) - GENERAL INFORMATION: Pass1 reconstruction analysis and four train user analysis currently in production
 - ◆ T0 site
 - ◇ GGUS:58718 [☞](#) submitted yesterday and concerning the bad performance of the CREAM-CE at CERN: SOLVED. The systems are back in production
 - ◇ GGUS:58587 [☞](#) concerning some missing c++ header files in 4 WN: waiting for news
 - ◆ T1 sites
 - ◇ All T1 sites in production
 - ◆ T2 sites
 - ◇ GGUS:58719 [☞](#) to GRIF_IRFU yesterday: CLOSED. (site in downtime)
 - ◇ GGUS:58750 [☞](#) to ru-PNPI today: local CREAM-CE not performing well. Site out of production
- LHCb reports -
 - ◆ T0 site issues:
 - ◇ xroot protocol adopted for CERN analysis jobs cured the problem reported previous days.
 - ◆ T1 site issues:
 - ◇ pic GGUS:58743 [☞](#) for a issue with FTS. All of attempts to submit FTS jobs to fts.pic.es were timing out with no response from the server [Gonzalo - can't comment on ticket yet, will follow up. Roberto - ticket solved!]

Sites / Services round table:

- INFN - ntr
- FNAL - ntr
- ASGC - ntr
- IN2P3 - two things: 1) transfers for ATLAS. Singled to LCG Fr technical director! Corresponding teams working on it. (FTS & network teams at T1). 2) Reminder than there will be AT RISK on June 8 - partial outage for MSS. Oracle DB patch to be applied - also on 8th. Should be transparent but have had bad experiences with this.
- NL-T1 - ntr
- PIC - ntr
- RAL - yesterday scheduled upgrade on DBs behind LFC and FTS for April CPU patch. Ran into problems and had to backout. Led to outage on these services for a couple of hours in afternoon.
- OSG - ntr
- GridPP - ntr

- CERN DB - one node of LHCb integration DB down today due to problem with f/s. For unknown reason got in R/O node. Problem fixed now.

AOB:

Friday

Attendance: local(Lola, Stephen, Simone, Jamie, Maria, Maarten, Roberto, Jean-Philippe, Manuel, Harry, Eva, Nilo);remote(Gang, Cristina, Jon, Doris, Michael, Rob, Joel, Tiju, Gonzalo, Rolf).

Experiments round table:

- ATLAS reports -
 - ◆ Today some instabilities have been observed in the ATLAS Central Catalog. The problem was related to the fact that Oracle decided to change the execution plan of a very popular query (the query itself is unchanged since many months). Oracle has gone to use HASH JOIN instead of NESTED LOOP and the HASH join needs a lot of memory. A intervention was scheduled at 14:00 to invalidate the Oracle cursor related to this query, involving a short outage of the CC since the httpd frontend needed to be stopped ;to flush the shared pool of instance 5 of the ATLR (this is the instance where DDM runs). But eventually, the undesired execution plan of the query has disappeared from the Oracle's shared pool memory without stopping the CC. Situation back to normal.
 - ◆ This morning it was decided to replicate all ESDs from the last reprocessing (150TB) to Nikhef. There are currently 5 copied of the ESDs worldwide (compare with 20+ copies of AODs) despite ESDs are the most popular format at the moment.
- CMS reports -
 - ◆ T0 Highlights
 - ◇ Processing normally
 - ◆ T1 Highlights
 - ◇ New skim passes starting
 - ◇ PromptSkimming resumed
 - ◇ PIC worried we were waiting on something from them due to new GGUS:58726 [↗](#) ticket
 - GGUS ticket should have been opened at the same time as the Savannah:114800 [↗](#) ticket
 - Was our error that it came later, sorry
 - ◆ T2 Highlights
 - ◇ MC production as usual
 - ◆ Can still not open ALARM tickets, opened GGUS:58764 [↗](#) yesterday.
- ALICE reports [↗](#) -
 - ◆ T0 site
 - ◇ GGUS:58587 [↗](#) concerning some missing c++ header files in 4 WN: updated with further information. The std.err files have been sent to the experts
 - ◆ T1sites
 - ◇ CNAF: Green light of the experiment to dismissing the local CASTOR system confirmed yesterday after the T1 service meeting
 - ◆ T2 sites
 - ◇ GGUS:58750 [↗](#) to ru-PNPI submitted yesterday: SOLVED. Site back in production
 - ◇ Legnaro: Bad performance of the local PackMan services at both local VOBOXES: SOLVED. Site back in production
- LHCb reports - Reprocessing of old data .
 - ◆ GGUS (or RT) tickets:
 - ◇ T0: 0
 - ◇ T1: 4
 - ◇ T2: 2
 - ◆ T0 site issues:

- ◇ LFC (RW) is reporting of alarms
 - Load on R/W instance of LFC - JPB in contact with Andy to find real cause of this problem; JPB - similar to what was seen on R/O instance a few weeks ago. A set of sessions opened and never closed. Given concrete examples with timestamped filenames to LHCb developers so that they can check.
- ◇ ~400 jobs failing at CERN: ~8 consuming too much memory. For others no mail received. Remedy ticket to LSF support.Manuel - working on it.
- ◆ T1 site issues:
 - ◇ CNAF Persistency access issue against Oracle GGUS:58770 [↗](#). Solution: Due to a problem on the private network interface of one of the two machines of the LHCb Oracle cluster, clients were losing the connection to the Conditions database at CNAF. Fixed.
 - ◇ CNAF StoRM: user reporting problems accessing data on GPFS GGUS:58794 [↗](#)
 - ◇ IN2p3 Persistency access issue against Oracle GGUS: 58766. It does not seem related to the PSU (IN2p3 did not apply it) but the problem concerns the memory allocation in the Shared Global Area [SGA] which is too small [1.5GB]. Increased.
 - ◇ PIC: CREAMCE, anomalous amount of jobs failing GGUS:58797 [↗](#).
- ◆ T2 sites issues:
 - ◇ INFN-NAPOLI-CMS and GRISU-CYBERSAR-CAGLIARI jobs aborting.

Sites / Services round table:

- FNAL - upgraded T3 and test instances of FTS to 2.2.4 and went fine. Will probably do T1 version next week.
- KIT - ntr
- IN2P3 - ntr
- BNL - made an interesting observation yesterday regarding storage servers. Running Solaris. Under heavy load due to TCP connections: then UDP packets not sent. Name resolution - primarily UDP based - suffering. Jobs not getting files staged as result. Got response and workaround from Sun overnight.
- CNAF - ntr
- ASGC - ntr
- RAL - ntr
- PIC - regarding CREAM CE issue reported by LHCb - already solved this morning. Problem detected to be in recent upgrade in Torque server. Versions of Torque > 2.4.6 there is a bug that means CREAM client not working properly. Had to roll back to 2.4.5 and working fine. CMS: Savannah/GGUS issue - any news? Stephen - working on it but no news. Reassign to dataops.
- OSG - several more tickets in q this morning, being worked on in a timely fashion - just an increase in ticket activity.

- CERN DB - repeat recommendation from CERN to Tier1 sites regarding Oracle April PSU: for those sites that have applied patch and have auditing & using COOL or similar (multiple client connections to same server) advised to roll-back patch.

AOB:

- Revised LHC schedule [↗](#)

-- JamieShiers - 28-May-2010

This topic: LCG > WLCGDailyMeetingsWeek100531

Topic revision: r13 - 2010-06-11 - PeterJones



Copyright &© 2008-2021 by the contributing authors. All material on this collaboration platform is the property of the contributing authors.
or Ideas, requests, problems regarding TWiki? use [Discourse](#) or [Send feedback](#)